




*FIG. 1 – Data Network with Edge Server*



**FIG. 2 – Service-Level Controlled Server**



***FIG 3 – Six Node Cluster***



*Fig. 4 – Service-Level Flow for Unmodified Applications*




Fig. 5 – Service-Level Flow for Modified Applications




Fig. 6 – Service-Level Flow for Unmodified Applications

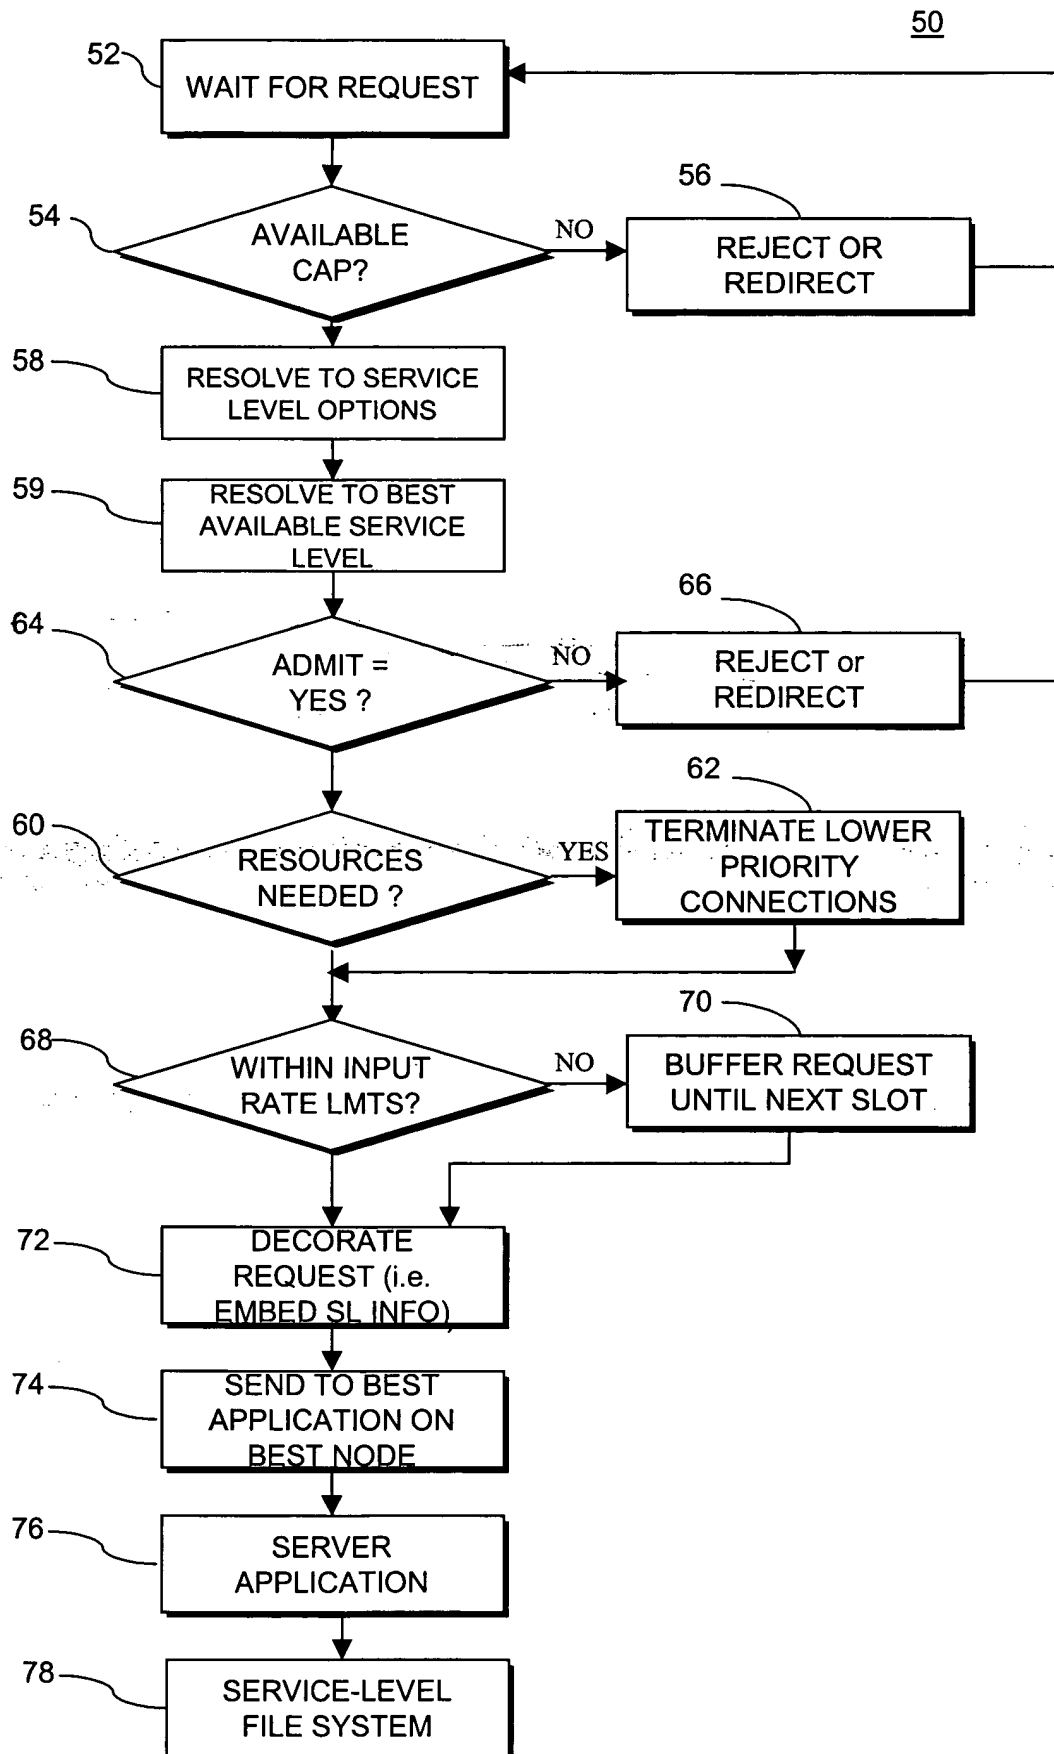
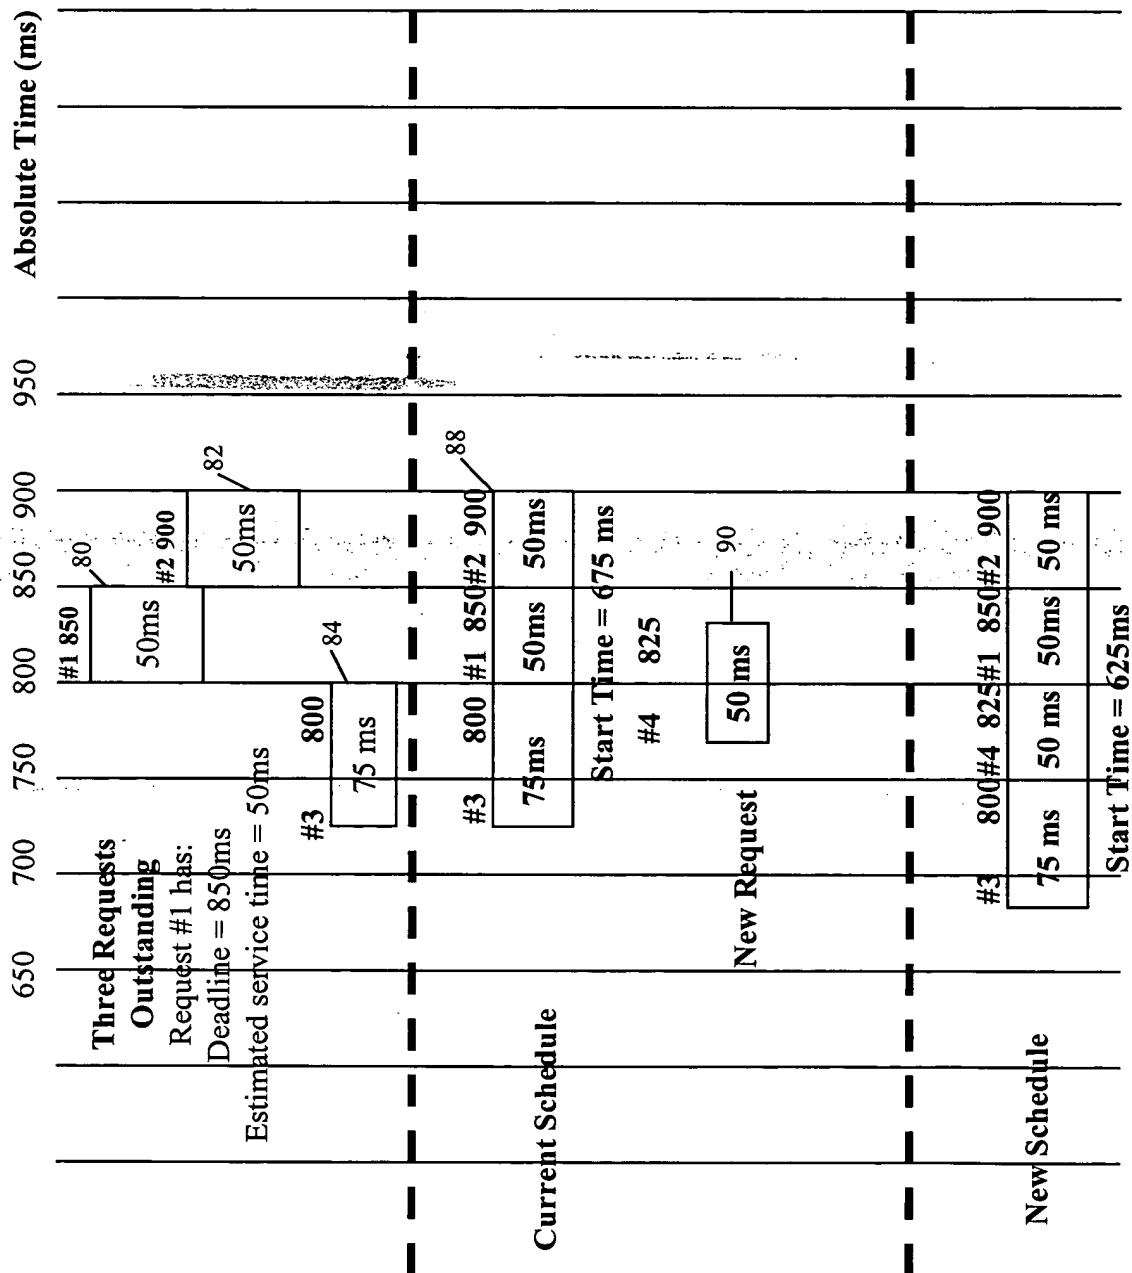


FIG. 7 = Service-Level Manager Flow



*Fig. 8 - Managing a Disk Schedule with a Single Priority*



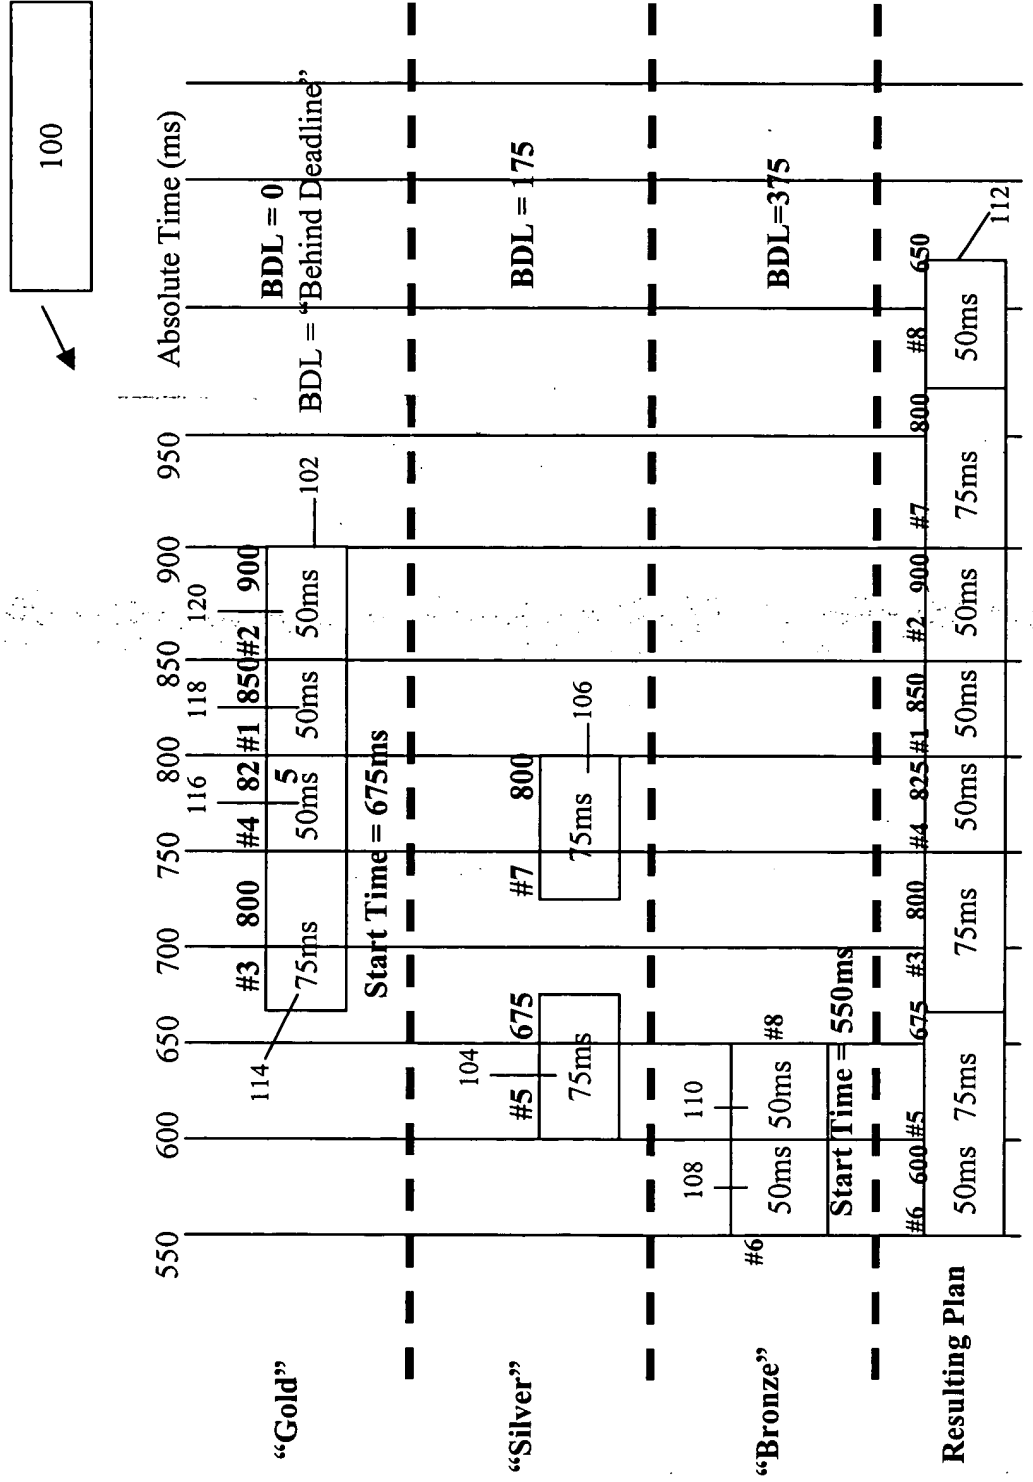
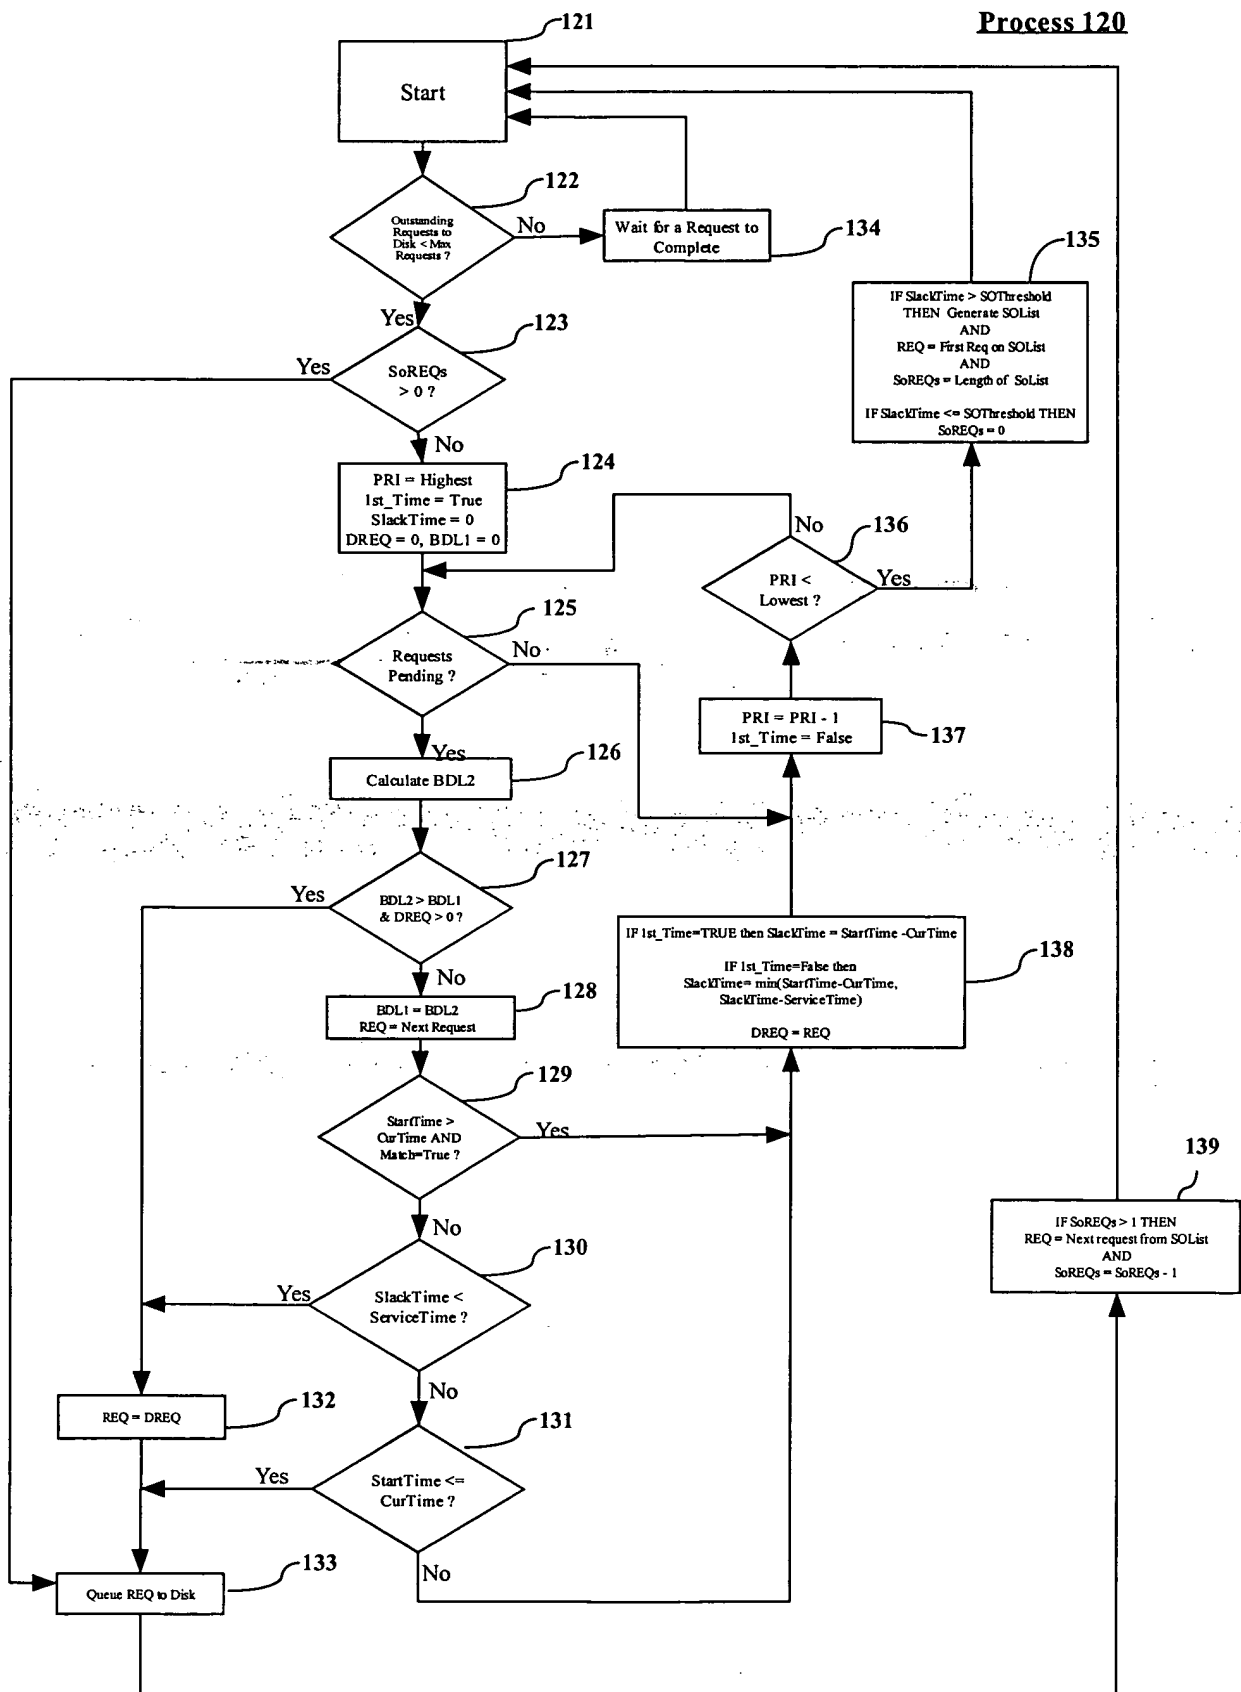


Fig. 9 - Managing a Disk Schedule with a Three Priorities



*Fig. 10 – Service-Level Disk Scheduler*

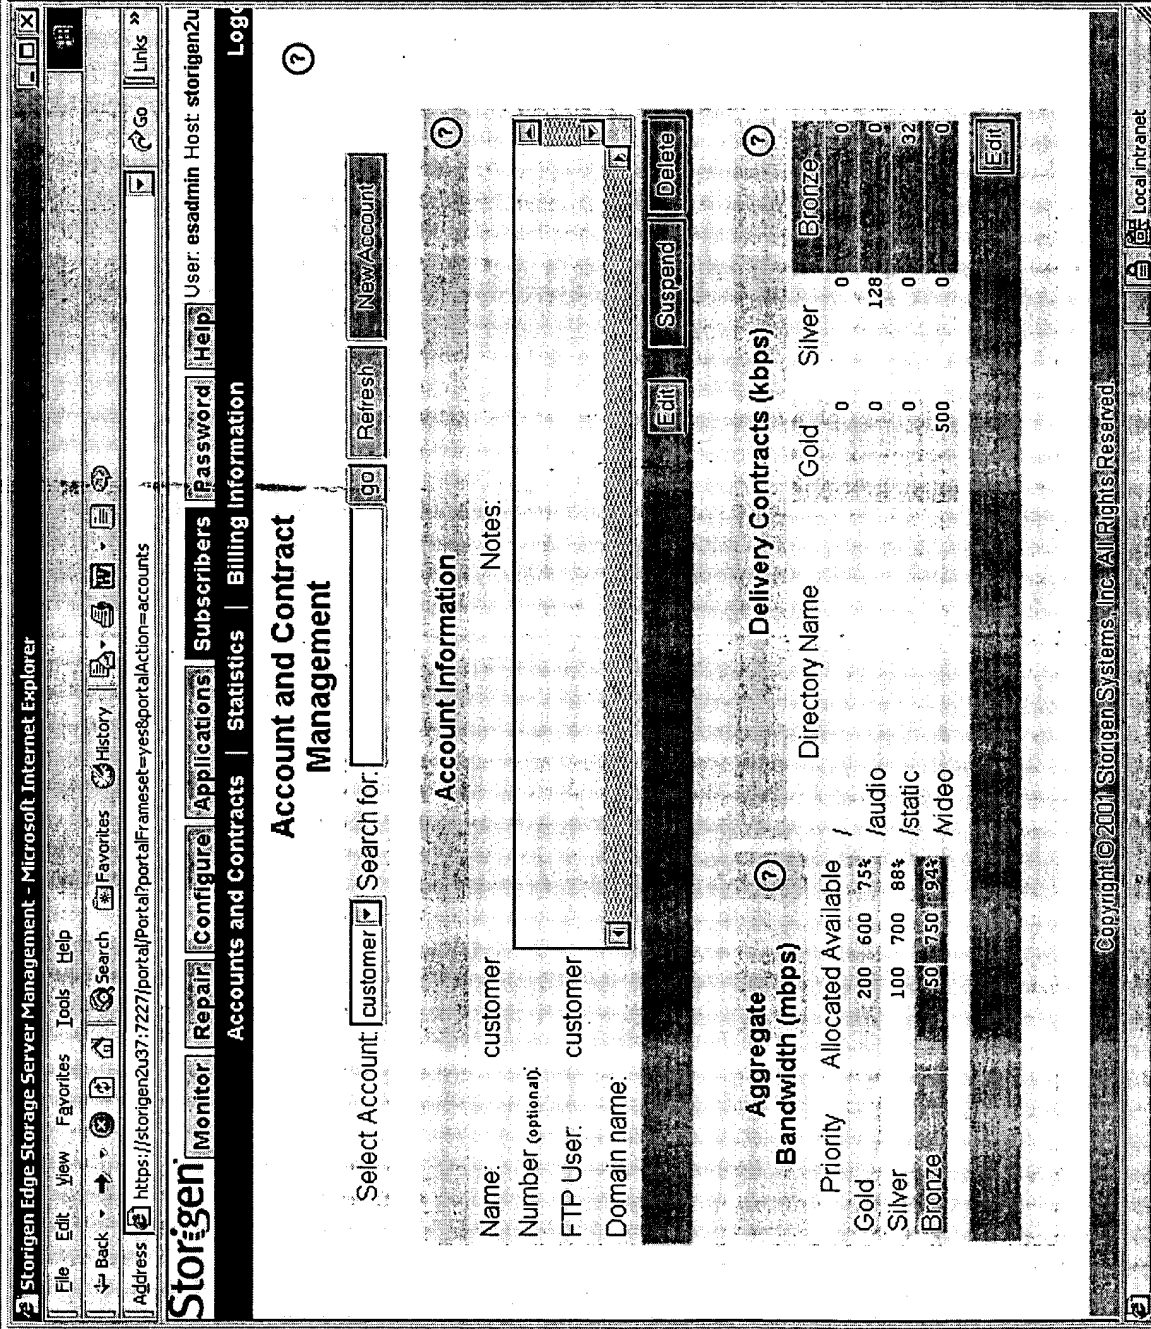



FIG. 11 – Example of Service-Level Management Interface

DIRECTORY



*FIG. 12 – Example of Directory Hierarchy for A Customer Account*

210

Based on request attributes generate inputs  
(Content/Requester/Application group identification)  
required for Service Level Resolver

212

Resolve Inputs to a set of Service Level options  
based on  
(Content/Requester/Application group identification)  
Service Level Tables

214

Resolve Service Level options to the best available  
Service Level based on committed capacity

218

*FIG. 13 – Process to Resolve Service Levels*

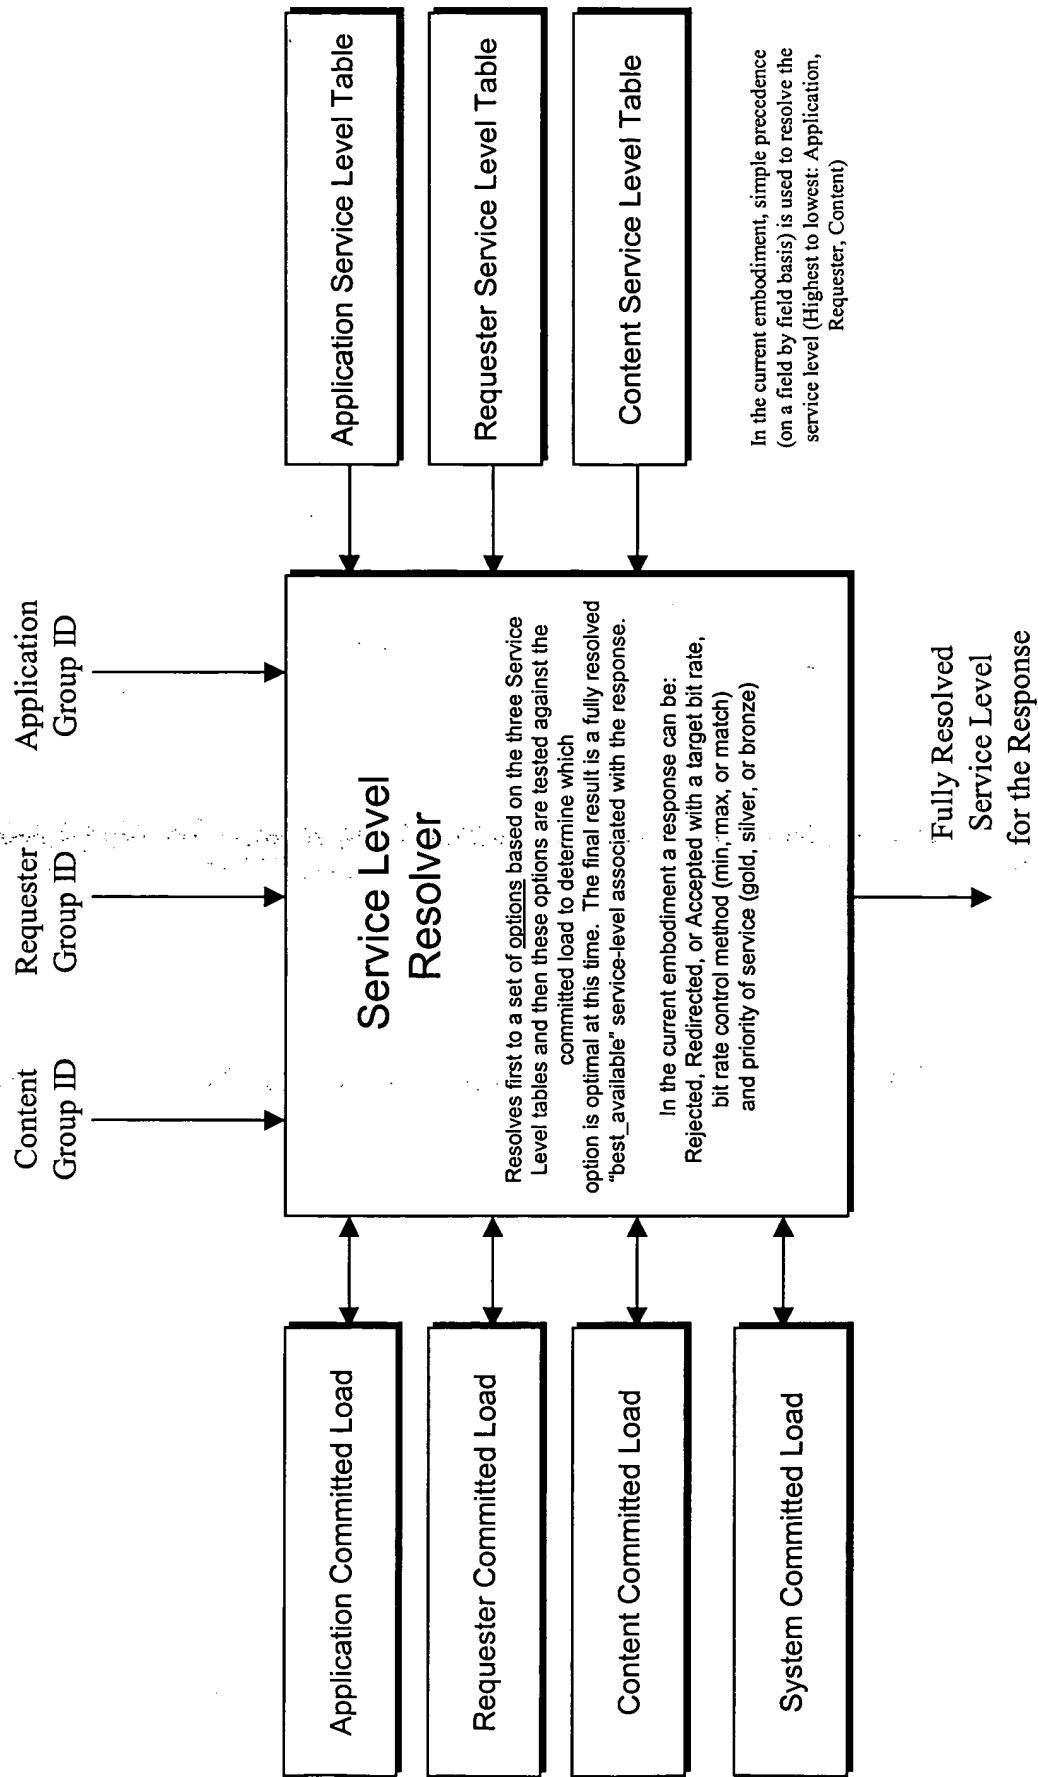


FIG. 14 - Service-Level Resolver

xGID	Priority	Aggregate Bandwidth (Mbps)	Individual Response Bit Rate (Kbps)	Bit Rate Control	On “No Admission”	URL for Redirection and Network QOS Settings
0001	GOLD	100	1.000	Min	Redirect	<a href="http://redirect1.com">http://redirect1.com</a>
0001	SILVER	100	0.500	Min	Reject	-
0001	BRONZE	200	0.100	Max	Reject	-
0002	GOLD	50	1.000	Min	Redirect	<a href="http://redirect2.com">http://redirect2.com</a>
0002	SILVER	0	0.000	Min	Reject	-
0002	BRONZE	0	0.000	Min	Reject	-

*FIG. 15 - Service Level Table Structure*

Priority Number (PN)	Service-level "Attribute"	"Color-based" Priority within a given Service-Level "Source"
10	Application	GOLD
9	Requester	GOLD
8	Content	GOLD
7	Application	SILVER
6	Requester	SILVER
5	Content	SILVER
4	Application	BRONZE
3	Requester	BRONZE
2	Content	BRONZE

*In the current embodiment, 1 = lowest priority.*

*A given priority level in general is not oversubscribed, but may allocate the entire capacity of the system.  
The system is over-subscribed as a result of the aggregate capacity allocated by all priority levels.*

**FIG. 16 - Priority Assignment Example**



Application GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0001	-	-	-	-	-
0001	-	-	-	-	-

Requester GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0001	-	-	-	-	-
0001	-	-	-	-	-

Content GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0004	8	100	1.000	Min	Redirect
0004	5	50	0.500	Min	Reject

In this example, the content-based service levels comprise the response options:

Service-Level Options	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
Option #1	8	100	1.000	Min	Redirect
Option #2	5	50	0.500	Min	Reject

There is more than 1Mbps spare capacity at Priority 8; therefore the response will be:

Best Available Service-Level	Priority	Response Bit Rate	Bit Rate Control	Action
Option #1	8	1.000	Min	Admit

In the current embodiment, -1 in a field indicates no entry (marked '-' above)

FIG. 17 - Content-based SL Example

Application GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0001	-	-	-	-	-
0001	-	-	-	-	-

Requester GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0001	9	6	2.000	Min	Redirect
0001	6	2	0.500	Min	Reject

Content GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0004	-	-	-	-	-
0004	-	-	-	-	-

In this example, the request-based service levels comprise the response options:

Response Options	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
Option #1	9	6	2.000	Min	Redirect
Option #2	6	2	0.500	Min	Reject

In this example, there is more than 2Mbps spare capacity at Priority 9; therefore the response will be:

Resolved Response	Priority	Response Bit Rate	Bit Rate Control	Action
Option #1	9	2.000	Min	Admit

FIG. 18 - Requester-based SL Example

Application GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0001	10	100	1.000	Min	Redirect
0001	7	50	0.500	Min	Reject

Requester GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0001	-	-	-	-	-
0001	-	-	-	-	-

Content GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0004	-	-	-	-	-
0004	-	-	-	-	-

In this example, the content-based service levels comprise the response options:

Response-Options	Priority	Aggregate-BW	Response-Bit-Rate	Bit-Rate-Control	On-"No-Admission"
Option-#1	10	100	1.000	Min	Redirect
Option-#2	7	50	0.500	Min	Reject

There is no capacity at priority 10 for this request but there is more than 0.5 Mbps spare capacity at Priority 7; therefore the response will be:

Resolved-Response	Priority	Response-Bit-Rate	Bit-Rate-Control	Action
Option-#1	7	0.500	Min	Admit

FIG. 19 - Application-based SL Example

## Lower Priority Application Service overrides Standard Content Service

Application GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0001	4	75	-	-	-
0001	-	-	-	-	-

Requester GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0001	-	-	-	-	-
0001	-	-	-	-	-

Content GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0004	8	100	1.000	Min	Redirect
0004	5	50	0.500	Min	Reject

In this example, the application-based service level defines the priority (it has precedence...). Since the application-based priority is lower than the priorities of the other potential service levels, there will be only one option. The content-based service level for the priority closest to the app-based priority defines the rest of the option:

Response:Options	Priority	Aggregate BW	Response:Bit:Rate	Bit:Rate:Control	On "No:Admission"
Option #1	4	75	0.500	Min	Reject
Option #2	-	-	-	-	-

There is more than 0.5 Mbps spare capacity at Priority 4; therefore the response will be:

Resolved:Response	Priority	Response:Bit:Rate	Bit:Rate:Control	Action
Option #1	4	0.500	Min	Admit

FIG. 20 - Application+Content-based SL Example

## Premium Requester Service overrides Standard Content Service

Application GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0001	-	-	-	-	-
0001	-	-	-	-	-

Requester GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0001	9	1	-	-	-
0001	-	-	-	-	-

Content GID	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
0004	8	100	1.000	Min	Redirect
0004	5	50	0.500	Min	Reject to www.redirect1.com

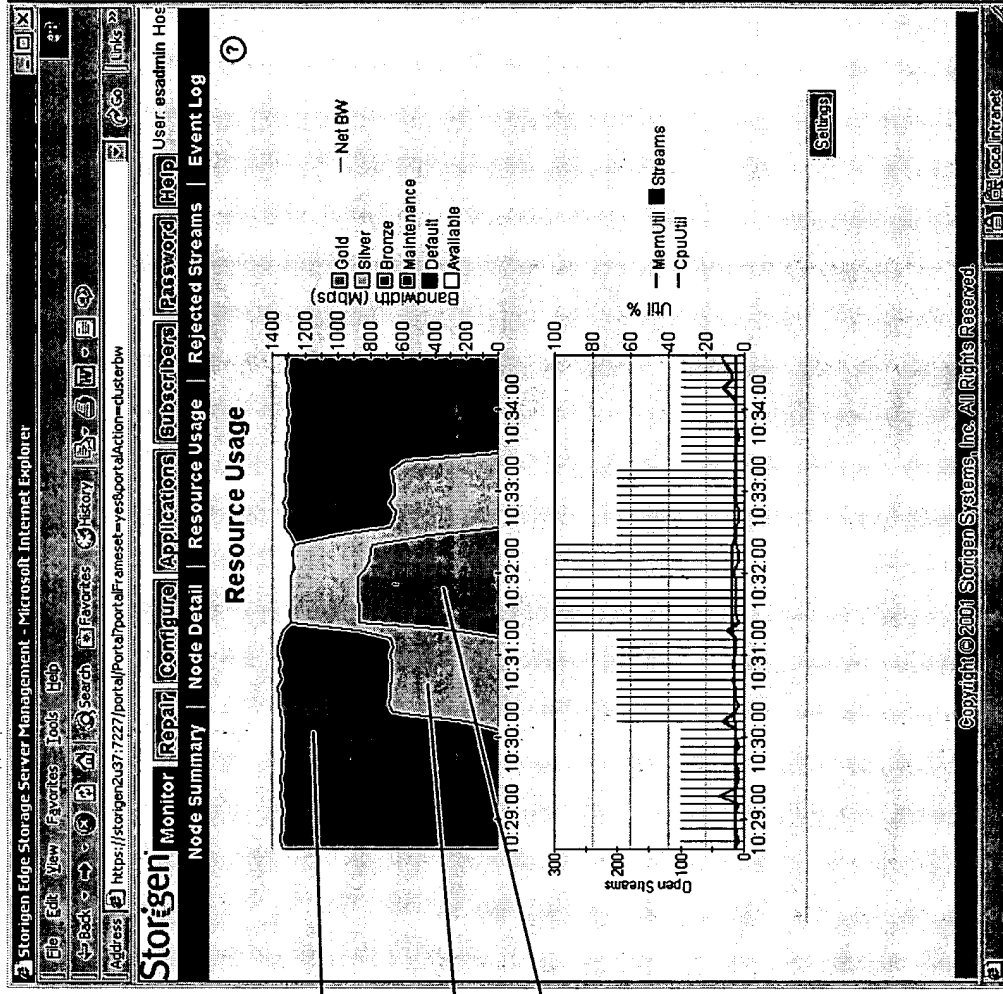
In this example, the requester-based service level defines the priority of Option#1  
The content-based service level for the priority closest to the app-based priority defines the rest of Option#1:  
Option 2 and 3 are taken directly from the Content Service Level entries.

Response Options	Priority	Aggregate BW	Response Bit Rate	Bit Rate Control	On "No Admission"
Option#1	9	1	1.000	Min	Reject to www.redirect1.com
Option#2	8	100	1.000	Min	Redirect
Option#3	5	50	0.500	Min	Reject to www.redirect1.com

There is 1 Mbps spare capacity at Priority 9; therefore the response will be:

Resolved Response	Priority	Response Bit Rate	Bit Rate Control	Action
Option#1	9	1.000	Min	Admit

FIG. 21 - Requester+Content-based SL Example



Bronze: 100 Users

Silver: 100 Users

Gold: 100 Users

FIG. 22 – Mixed Service Level Workload